The DOE ACTS Collection

The DOE Advanced CompuTational Software Collection

The Department of Energy's Advanced CompuTational Software (ACTS) Collection is a set of DOE-developed software tools that make it easier for programmers to write high-performance scientific applications for parallel computers. The ACTS Collection is maintained by the Scientific Computing Group in Berkeley Lab's Computational Research Division with support from the Mathematical, Information, and Computational Sciences (MICS) Division of DOE's Office of Science.

During the past several decades there has been continuous growth in the number of physical and societal problems that have been successfully studied and solved by means of computational modeling and simulation. A number of these are important scientific problems ranging in scale from the atomic to the cosmic. For example, ionization is a phenomenon as ubiquitous in modern society as the glow of fluorescent lights and the etching on silicon computer chips; but it was not until 1999 that researchers finally achieved a complete numerical solution to the simplest example of ionization, the collision of a hydrogen atom with an electron. On the opposite scale, cosmologists have long wondered whether the expansion of the universe, which began with the Big Bang, would ever reverse itself, ending the universe in a Big Crunch. In 2000, analysis of new measurements of the cosmic microwave background radiation showed that the geometry of the universe is flat, and thus the universe will continue expanding forever.

Both of these discoveries depended on high-performance computer simulations that utilized computational tools included in the ACTS Collection. The ACTS Collection is an umbrella project that brings together a number of general-purpose computational tool development projects funded and supported by DOE. These tools, which have been developed independently, mainly at DOE laboratories, make it easier for scientific code developers to write high performance applications for parallel computers. They tackle a number of computational issues that are common to a large number of scientific applications, mainly implementation of numerical algorithms, and support for code development, execution and optimization. The ACTS Collection Project enables the use of these tools by a much wider community of computational scientists, and promotes code portability, reusability, reduction of duplicate efforts, and tool maturity.

The tools fall into these broad categories:

NUMERICS: The numerical tools implement numerical methods widely used in the solution of PDEs, ODEs and optimization problems found in many large scientific applications.

CODE DEVELOPMENT: These tools provide infrastructure that manages some of the complexity of parallel programming (e.g., distributing arrays, communicating boundary information, etc.) but do not actually implement numerical methods.

EXECUTION SUPPORT: Execution support is a catch-all category for application-level tools; these tools include performance analysis and remote visualization support.

DEVELOPER SUPPORT: These tools provide a infrastructure for tool developers and probably will not be used or seen directly in scientific applications.

Here are some of the current tools in the ACTS Toolkit:

Numerical Tools

TRILINOS is an effort to bring modern object- oriented software design to high-performance parallel solver libraries for the solution of large-scale complex multi-physics engineering and scientific applications.

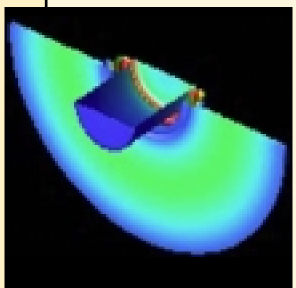

Computation of eigenfrequencies and eigenmodes of three-dimensional inhomogeneous resonant cavities using ARPACK and PETSc.

HYPRE is a library of linear system preconditioners which are designed both for stand alone use and for interoperability with other ACTS tools.

PETSC (Portable, Extensible Toolkit for Scientific Computation) provides many of the parallel algebraic data structures, solvers, and related infrastructure required to scalably solve PDEs using implicit methods on finite elements, finite differences, or finite volumes.

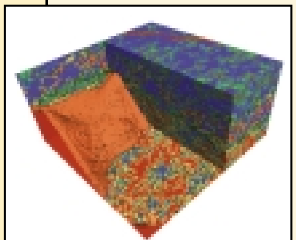

This figure is a 3D simulation of fluid flow through a heterogeneous porous media performed by the ParFlow Project using SUNDIALS.

SUNDIALS (SUite of Nonlinear and DIfferential/ALgebraic equation Solvers), consists of ODE solvers, DAE solvers, and nonlinear system solvers systems.

The **ScaLAPACK** library extends LAPACK's high-performance linear algebram software to distributed memory, message passing MIMD computers and networks of workstations supporting PVM or MPI.

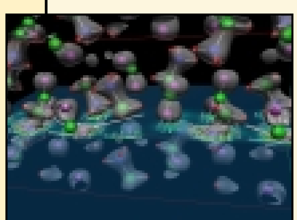

Material Science Application that uses ScaLAPACK. This image shows the induced current (white arrows) and charge density (colored plane and gray surface) in crystallized glycine due to an external field.

SUPERLU is a general purpose library for the direct solution of large, sparse, nonsymmetric systems of linear equations on high performance machines. Serial, shared memory and distributed memory implementations are available.

TAO The Toolkit for Advanced Optimization (TAO) focuses on large-scale optimization software, including nonlinear least squares, unconstrained minimization, bound constrained optimization, and general nonlinear optimization.

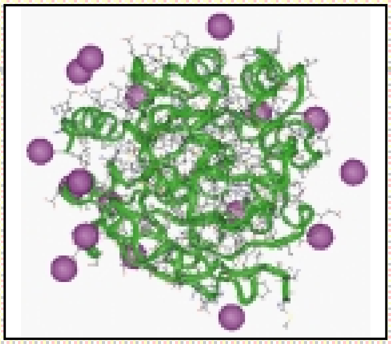

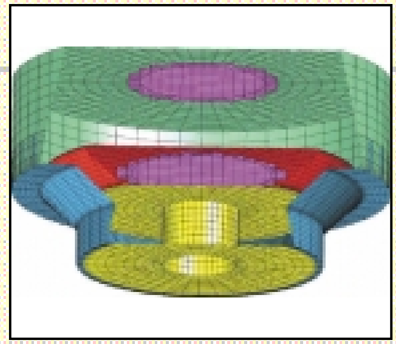

Tools Facilitating Code Development

GLOBAL ARRAYS is a portable, distributed array library which provides a shared memory style of programming without hiding the nonuniform access characteristics of the arrays.

OVERTURE is a library of grid functions which derives (formally) from P++ arrays and provides vertex or cell centered grids. It includes support for complex geometry and graphical visualization.

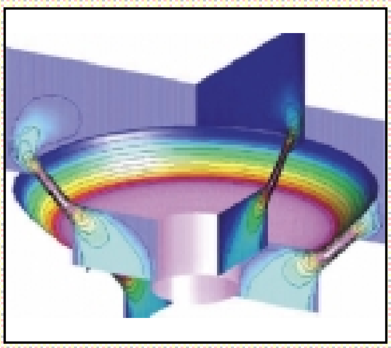

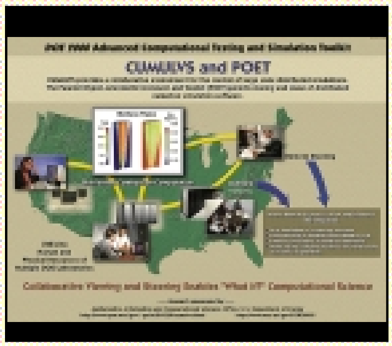

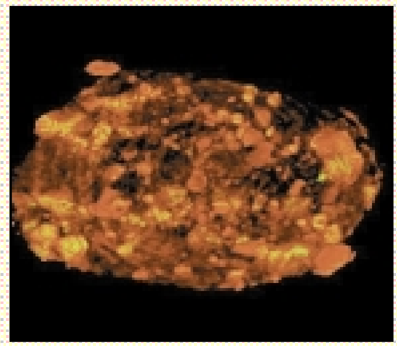

Tools Supporting Code Execution and Development

CUMULVS (Collaborative User Migration User Library for Visualization and Steering) provides a mechanism which enables programmers to incorporate computational steering, interactive visualization, and fault tolerance into existing parallel programs.

GLOBUS provides a means for the creation of large scale Computational Grids and a toolkit of core services used by applications to access the grid.

PAWS (Parallel Application WorkSpace) provides application support and inter-application data transfer capabilities in heterogeneous distributed environments. Features include application launching, computational steering, and transfer of distributed data.

SILOON (Scripting Interface Languages for Object-Oriented Numerics) is destined to aid in the writing of the scripts that can be used in some runtime systems such as PAWS.

TAU (Tuning and Analysis Utilities) is a set of tools that allow you to analyze the performance of distributed and multithreaded programs. It is especially useful in tracking C++ classes that are created and destroyed dynamically.

ATLAS (Automatically Tuned Linear Algebra Software) and PHiPAC (Portable High Performance ANSI C) are tools for the automatic generation of optimized numerical software for modern computer architectures. They can match hand-tuned performance for level three BLAS operations.